that particular person. Any walk whose likelihood given a distribution is greater than a threshold is assumed to be from the person that the distribution represents, and is used in computing gait parameter estimates for that person. The classification may be performed independently for each distribution. Thus, a walk could be included in the estimates of more than one person, if the distributions overlap. The steps of model initialization and updating are described below and illustrated in FIG. **2**.

An output module **320** processes the depth image data **112** and/or the one or more generated parameters **114** to perform one or more health risk assessments. For example, the parameters **114** and depth image data **112** may be used to assess a patient's risk of falling and/or the onset of illness.

In one embodiment, the actual assessment of fall/health risk may be based on mapping the various gait parameters to standard clinical measures such as a Timed-up-and-Go (TUG) test, and the Habitual Gait Speed (HGS) test. For example, in one embodiment, a simple neural network model that can "predict" TUG time based on an individual person's average gait speed. It is contemplated that any gait parameter may be mapped to standard measures. For example, a TUG time above 16 or 20 seconds indicates a high risk of falling in the next year. Accordingly, the gait parameter data and/or gait parameter estimates may be used to predict a score that a person, such as a patient, would receive on various clinical measures, tests, and the like, such as the TUG, HGS, Berg Balance-Short Form, Short Physical Performance Battery, and Multi-Directional Reach Test data, etc.

FIG. **4** depicts an example method and/or process **400** for obtaining depth image data and subsequently processing the depth image data to generate temporal and spatial gait parameters for use in health risk assessment. Process **400** may be executed by at least one processor encoded with, or executing instructions of, an image analysis application **109**. Initially, at **402**, process **400** includes receiving depth image data for a particular patient from one or more depth cameras **108**. For example, depth image data may be received from a Microsoft Kinect™ camera device located in the home of an elderly gentleman. At **404**, the depth image data may be analyzed to generate at least one three-dimensional object. For example, a three-dimensional representation of the elderly gentlemen patient may be generated. At **406**, a walking sequence may be identified based on the at least one three-dimensional object. For example, a walking sequence corresponding to the elderly gentleman patient may be identified. At **408**, one or more parameters may be generated from the walking sequence. For example, one or more temporal and spatial gait parameters may be generated corresponding to the elderly gentleman. At **410**, the generated parameters are used to perform various health risk assessments for a particular patient (e.g. the elderly gentlemen) and results of the health risk assessments may be provided for display at **412**.

The description above includes example systems, methods, techniques, instruction sequences, and/or computer program products that embody techniques of the present disclosure. However, it is understood that the described disclosure may be practiced without these specific details. In the present disclosure, the methods disclosed may be implemented as sets of instructions or software readable by a device. Further, it is understood that the specific order or hierarchy of steps in the methods disclosed are instances of example approaches. Based upon design preferences, it is understood that the specific order or hierarchy of steps in the method can be rearranged while remaining within the disclosed subject matter. The accompanying method claims present elements of the various steps in a sample order, and are not necessarily meant to be limited to the specific order or hierarchy presented.

The described disclosure may be provided as a computer program product, or software, that may include a machine-readable medium having stored thereon instructions which may be used to program a computer system (or other electronic devices) to perform a process according to the present disclosure. A machine-readable medium includes any mechanism for storing information in a form (e.g., software, processing application) readable by a machine (e.g., a computer). The machine-readable medium may include, but is not limited to, magnetic storage medium (e.g., floppy diskette); optical storage medium (e.g., CD-ROM); magneto-optical storage medium; read only memory (ROM); random access memory (RAM); erasable programmable memory (e.g., EPROM and EEPROM); flash memory; or other types of medium suitable for storing electronic instructions.

It is believed that the present disclosure and many of its attendant advantages will be understood by the foregoing description, and it will be apparent that various changes may be made in the form, construction and arrangement of the components without departing from the disclosed subject matter or without sacrificing all of its material advantages. The form described is merely explanatory, and it is the intention of the following claims to encompass and include such changes.

While the present disclosure has been described with reference to various exemplary embodiments, it will be understood that these embodiments are illustrative and that the scope of the disclosure is not limited to them. Many variations, modifications, additions, and improvements are possible. More generally, embodiments in accordance with the present disclosure have been described in the context of exemplary implementations. Functionality may be separated or combined in blocks differently in various embodiments of the disclosure or described with different terminology. These and other variations, modifications, additions, and improvements may fall within the scope of the disclosure as defined in the claims that follow.

What is claimed is:

1. A method comprising:

receiving, by at least one processor, depth image data from at least one depth camera, wherein the depth image data comprises a plurality of frames that depict a person walking through a home environment over time, the frames comprising a plurality of pixels;

performing, by the at least one processor, segmentation on the pixels of the frames;

in response to the segmentation, (1) generating, by the at least one processor, a three-dimensional (3D) data object based on the depth image data, and (2) tracking, by the at least one processor, the 3D data object over a plurality of frames of the depth image data, wherein the tracked 3D data object comprises time-indexed spatial data that represents the person walking through the home environment over time;

identifying, by the at least one processor, a walking sequence from the tracked 3D data object, wherein the identifying step comprises:

the at least one processor determining a speed for the tracked 3D data object over a time frame;

the at least one processor comparing the determined speed with a speed threshold;

in response to the comparison indicating that the determined speed is greater than the speed threshold, the at least one processor assigning a state indicative of walking to the tracked 3D data object;